March-April 1995
Vol. 2 No. 9

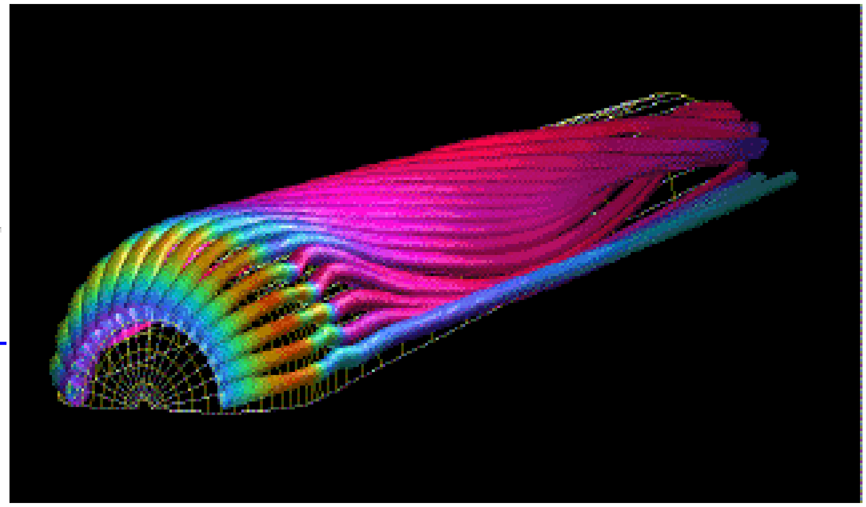

NAS-Supported Visualization Research Important for Interpreting Complex Tensor Data

by Kristina D. Miceli

Researchers at Stanford University, working under a NAS-sponsored visualization grant, are devising creative techniques to help scientists depict the complex information contained in tensor data. Principal investigator Lambertus Hesselink and graduate student Thierry Delmarcelle focused the first year of their research on hyperstreamlines and tensor field topology.

Few Tools for Viewing Tensor Data

Computer-generated visualizations of computed flow fields typically mimic the images obtained from experimental studies. For example, streamlines computed from computational fluid dynamics (CFD) simulations mimic the patterns created when smoke is injected in windtunnel studies. Techniques for visualizing scalar and vector field data are widely available and used by scientists. In contrast, techniques to visualize tensor fields -- fundamental in engineering and physical sciences -- are not readily available, making it difficult for scientists to effectively analyze tensor data. Reasons for the lack of visualization techniques for tensor data include the complex, multivariate nature of the data and the fact that no commonly used experimental analogy exists for visualizing tensor data. People are simply not used to creating mental images of tensor data.

Second-order tensor fields, which consist of a 3 x 3 matrix defined at each node in a computational grid, are frequently computed by researchers in computational aerosciences disciplines such as fluid dynamics and structural dynamics. Tensor field variables in CFD include stress, viscous stress, rate of strain, and momentum flux; tensor variables in structural dynamics include stress and strain.

Because of the lack of visualization techniques for analyzing all components of a complex tensor field, scientists typically resort to either simplifying the tensor field data and visualizing it as a scalar field or visualizing the individual vectors that comprise the tensor field. Both methods result in the loss of valuable information that might provide new insight into physical problems.

Hyperstreamlines Explained

Vector fields can be visualized effectively using streamlines, which represent the trajectory path of a particle in a flow field. Hyperstreamlines, an extension of the concept of streamlines to second-order tensor fields, provide the scientist with a continuous representation of the tensor field along a three-dimensional path.

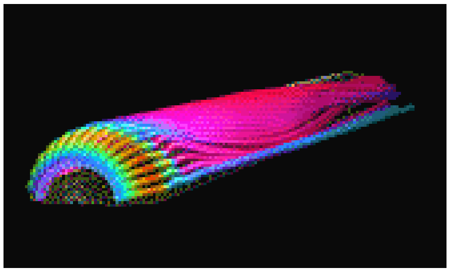

Hyperstreamlines are defined by a trajectory path and a cross-sectional geometry at each point in that path. The trajectory path of a hyperstreamline is essentially a streamline computed from one of the eigenvector fields. The cross-sectional geometry of a hyperstreamline is determined by the two remaining eigenvector fields and can either be an ellipse, a circle (if the eigenvectors are "degenerate," or equal to each other), or a cross. Connecting these cross-sectional geometries along the trajectory path results in a hyperstreamline that either looks like a tube (elliptical/circular geometry) or a helix (cross geometry). Figure 1 shows a collection of tube-shaped hyperstreamlines.

Topology Depicts Global Structure

Visualizing the topology of a tensor field is an effective way for scientists to display the complex structure of an entire tensor field, providing a more global representation of the data than hyperstreamlines, which only depict a subset of the data.

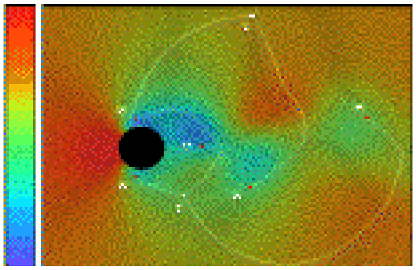

Tensor field topology involves a detailed mathematical analysis of a tensor data field. Through this analysis scientists can identify degenerate points in the flow field. Further analysis produces a topological skeleton that divides the tensor field into regions bounded by dividing lines that meet at the degenerate points. The resulting skeleton is a description of the structure of a flow field, indicating where transitions in the flow occur. Degenerate points and a topological skeleton are depicted in Figure 2.

Research Well Received

The research by Hesselink and Delmarcelle has received two prestigious awards in the visualization community: a paper describing this research, entitled "Visualizing Second Order Tensor Fields with Hyperstreamlines" was selected by IEEE Computer Graphics and Applications as the Best Paper for 1993; a conference paper, entitled "The Topology of Symmetric Second-order Tensor Fields" shared the best paper distinction (together with a paper presented by Lisa Forssell, a Stanford doctoral candidate and NAS researcher; *see NAS News, Jan-Feb 1995)* at Visualization '94 held last October in Washington, D.C.

Future Plans Focus on 3D Data

Hesselink will continue investigating tensor field visualization techniques during the second year of the research grant. [Delmarcelle recently received his doctorate based on this research and is now in Belgium]. Plans for future research include extending the topology research to three dimensions, and combining three-dimensional tensor data with three-dimensional vector data to examine possible benefits of correlating the information. To contact Hesselink, send email to: **bert@kaos.stanford.edu**.

Next Article | Contents | Main Menu | NAS Home

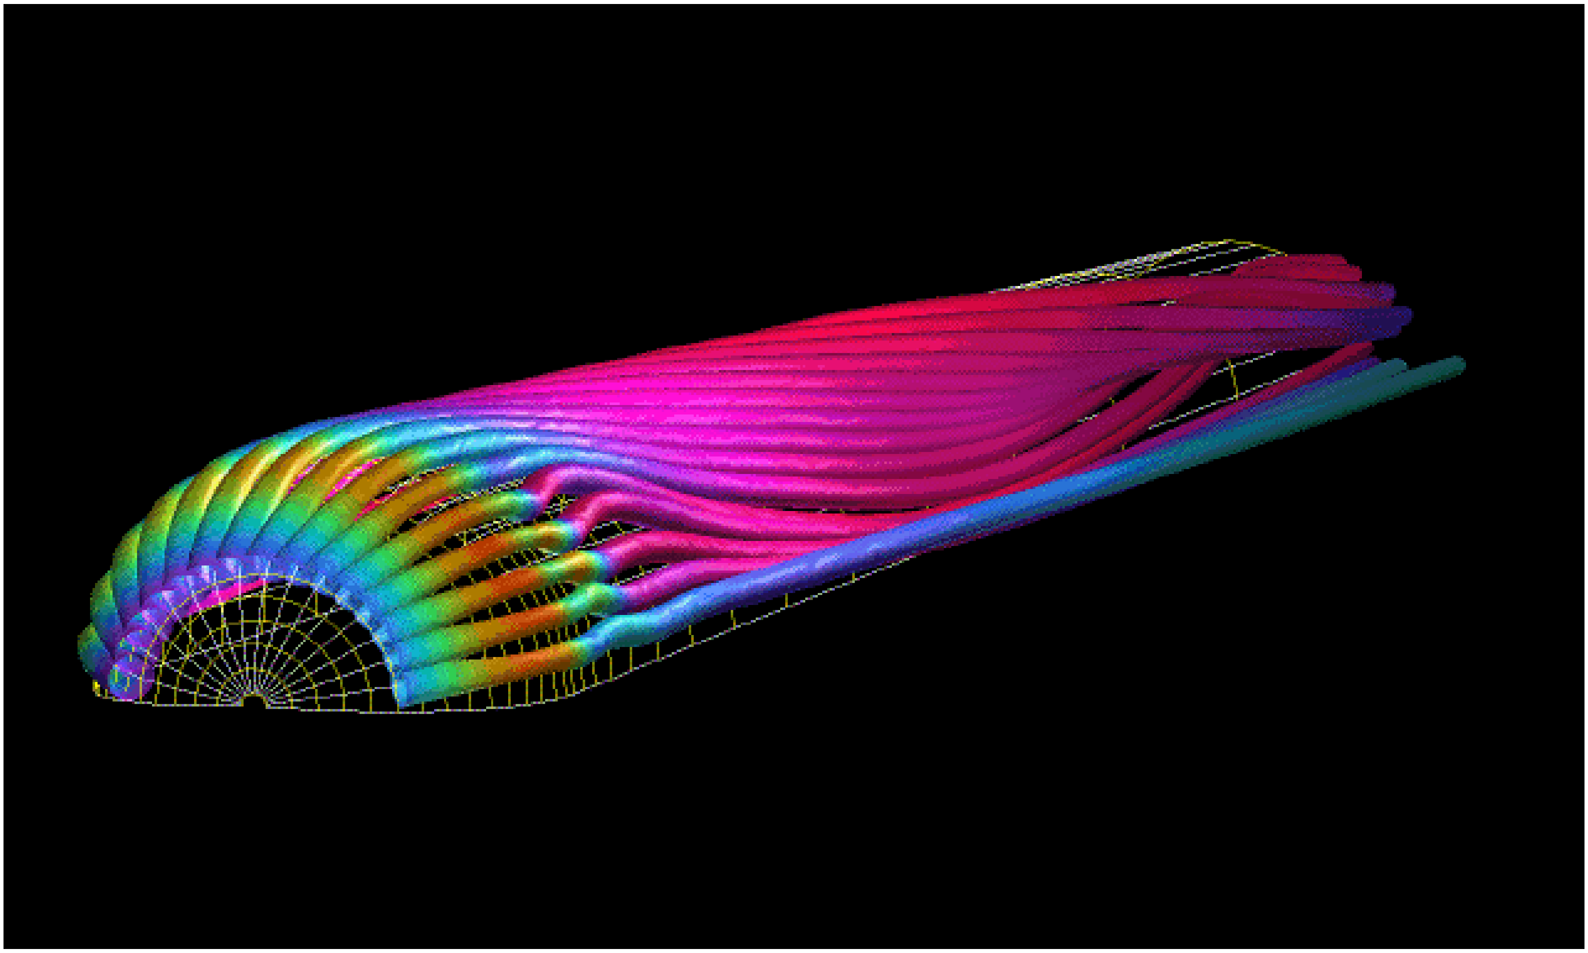

Figure 1: The reversible momentum flux density tensor depicted as hyperstreamlines in a flow past a hemisphere cylinder. The trajectories of the hyperstreamlines are streamlines of the velocity field. The eigenvalues used to compute the cross-sectional geometry are degenerate, resulting in a circular cross-section, with the diameter of the section encoded by local pressure. The color map represents the density of kinetic energy.

to the article

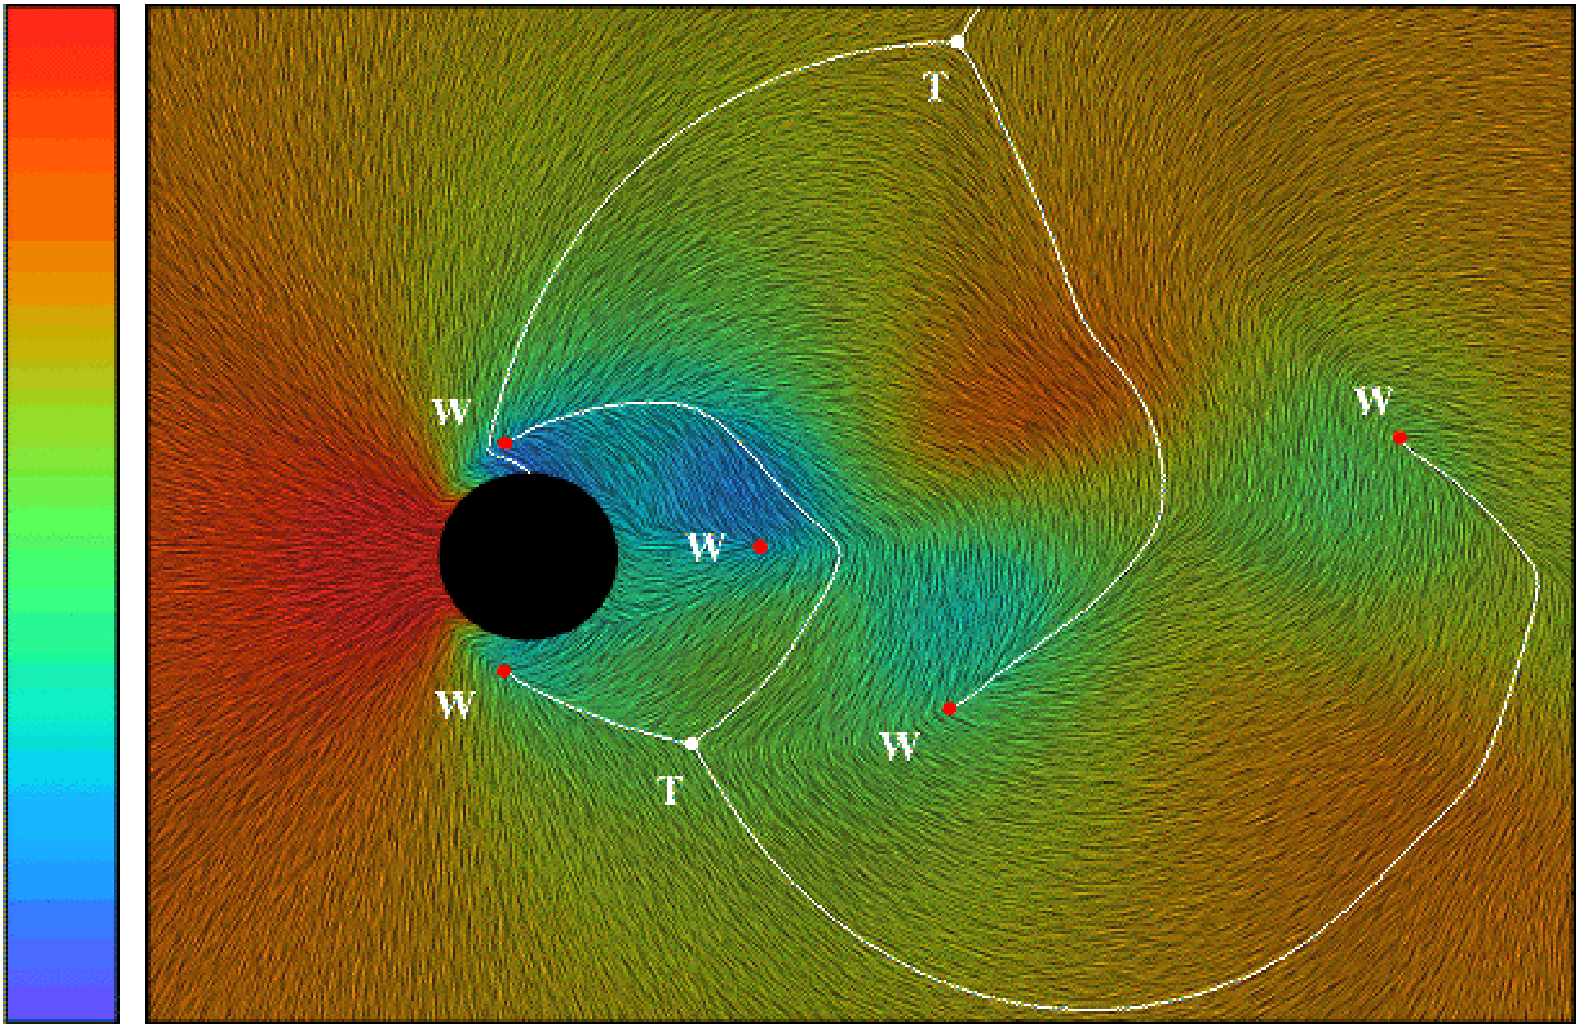

Figure 2: A two-dimensional topological skeleton representing the most compressive eigenvector of the stress tensor in a flow past a cylinder. The texture represents the same eigenvector field and is added for correlation purposes, since topological skeletons can be difficult to interpret by themselves. Color encodes the magnitude of the compressive force, from most compressive (red) to least compressive (blue). Wedge points (W) and trisector points (T) are also shown. to the article

NAS's 160-node SP2 `Scales Well'

by Elisabeth Wechsler

The 160-node IBM SP2 at NAS performed an average of 1.4 times faster than the 16-processor NAS CRAY C90 on three pseudo-application NAS Parallel Benchmarks (NPB) run on January 17. According to Toby Harness, of the NAS parallel systems group, "All three benchmarks did better than expected, which shows that the codes are scaling well," based on test results for the former 64-node and 128-node SP2 configurations. The tests were run on full-production NAS machines by NAS staff, Harness noted. A full report on the NPB, including these latest results, is available.

The 160-node configuration, which includes the final 32 nodes specified under the NASA Computational Aerosciences-funded Cooperative Research Agreement, allocates 144 nodes for scientific jobs during off hours (or 128 nodes during the day), reserving 16 nodes for system software development.

Other improvements to the SP2 include upgraded "general availability" IBM system software; six HiPPI channels in full production, allowing the SP2 to access NAS high-performance systems over the HiPPI network; and support for P-OVERFLOW, completed ahead of schedule.

P-OVERFLOW, a subset of OVERFLOW's most frequently used functionalities, is intended for distributed memory MIMD (multiple instruction, multiple data) computers, according to Sisira Weeratunga, who ported the code to the SP2. OVERFLOW is a general-purpose implicit flow solver, developed by Pieter Buning at Ames Research Center. It uses the concept of a composite of logically structured overset grids to solve flows around complex geometries.

The performance of P-OVERFLOW on the 160-node SP2 is "good," Weeratunga said, adding that "fewer than 8 nodes of the SP2 achieve the performance of 1 CPU of a C90."

Performance results comparing MPL, MPI, and PVMe for the 160-node system are also available.

Another SP2 project currently underway is the creation of a 2-node SP2 system with an Ethernet connection "to enable early software development to proceed without impacting the whole system," Harness said. "This will allow us to work out some problems before scaling up the software to the full 160-node SP2 system."

In addition, the front-end workstation for the SP2 was upgraded in January from an RS/6000 58H to an RS/6000 590, offering a faster processor and more disk space, according to Harness. A short-term future project is to convert Portland Group Inc.'s high-performance Fortran to the SP2 environment, he said.

Users Now Have 'Instant' Access to Beta Software From NAS

by Jill Dunbar

A waiver put in place by Lee Holcomb, director of NASA's High Performance Computing and Communications office -- along with a new procedure on NAS's World Wide Web (WWW) server -- now give users nearly instant access to a selection of NAS-developed beta-version software. The procedure supports the agency-wide decision to improve its effectiveness in transferring technology to industry.

The new method eliminates some annoying red-tape for nongovernment users, including a written request on company letterhead and the signing of a nondisclosure agreement.

Aside from saving time and hassles for users, this approach offers an opportunity for NAS and industry to collaborate in software development and test cycles, said David Yip, head of NAS's information management group. His group created the mechanism used to retrieve software from the WWW. Called the NAS Software Archive, it uses a larger database search system dubbed FERIT (Find, Explore, and Retrieve IT).

Cathy Schulbach, who heads the NAS tools group, advocates the quick-and-easy distribution method. "Ultimately, this creates a better product, and one that is more likely to be supported and commercialized by industry -- that's what technology transfer is all about." Schulbach's group developed AIMS, a toolkit that facilitates performance evaluation and tuning on parallel applications, and NTV, a tool for visualizing execution traces. Both programs are available through the archive system.

Other items in the NAS Software Archive include: MPIRUN, a multidisciplinary portable program loader; TOP 2.0, which displays information about operating systems and processes; and several computer security-related applications. The archive will continue to be expanded as development cycles progress.

Yip emphasized that some restrictions for NAS-distributed software remain: Access is limited to sites with mailing addresses within the U.S., and users must agree to not redistribute the software. If a company wants to make commercial use of NAS software, it must work with the NASA Ames Commercial Technology office. (For more information, contact, Syed Shariq, the office's director, at **syed_shariq@qmgate.arc.nasa.gov**.)

Completed Software Still via COSMIC

For the time being, users must continue to purchase production-ready software from COSMIC, NASA's Computer Software and Management Information Center. (COSMIC software is usually free to government users.) But, during his visit to NASA Ames Research Center last October, Goldin promised to investigate the issues surrounding COSMIC, after Ames center managers cited bottlenecks in using this distribution vehicle.

NAS-developed software available through COSMIC includes FAST (Flow Analysis Software Toolkit) -- one of COSMIC's most-requested items -- and HNMS 2.0 (Hierarchical Network Management System), designed to help network administrators manage large Internet Protocol networks.

Next Article | Contents | Main Menu | NAS Home

Ames Scientist Begins Astronaut Training

by Elisabeth Wechsler

Kalpana Chawla

Kalpana Chawla, an aerospace engineer working on CFD research at Ames Research Center, was selected by NASA in December as a Space Shuttle astronaut candidate for the 1995 class. She reports to Johnson Space Center (JSC), Houston, on March 6 to begin one year of training and evaluation as a future astronaut.

NASA brought 122 applicants, chosen from a pool of 2,962 applicants, to JSC for interviews and medical evaluations last summer. Six civilian and 13 military officer candidates were selected. Chawla is one of nine "mission specialists," joining a group of ten pilots. Three international candidates will be announced later, according to NASA officials.

Survival, Aircraft Operations

The first year of training introduces candidates to all NASA centers, instructs them in land and sea survival, and prepares them for operating aircraft. After successful completion of training, the candidates become "astronauts" and are assigned specific jobs/duties within the NASA Astronaut Office to further prepare them as shuttle crew members.

"It's more important to have a proven track record as a scientist rather than knowledge in any specific field," Chawla said. "Even though I have a technical degree in aerospace engineering, NASA is looking for a candidate's ability to solve technical problems. It's how someone tackles a problem that's important," she said.

The mission specialists represent a diverse group of scientific disciplines, including physics, aerospace engineering, space sciences, chemistry, computer science, geology, and astrophysics. Their typical functions will include operating the manipulator arm for retrieving and launching satellites, conducting and monitoring experiments, and performing EVAs (Extra-Vehicular Activities).

Ames Work `Particularly Useful'

Chawla believes that her exposure to a wide variety of computer systems at Ames will be particularly useful in her new role. As a research scientist and vice president of Overset Methods Inc., Los Altos, CA, she worked at Ames in the Applied Computational Aerodynamics Branch. One project involved unsteady powered-lift simulations of a delta wing in descent on the NAS CRAY Y-MP. These simulations resulted in very large datasets that were processed on the NAS Time Accurate Visualization System, a Convex C3240.

A typical space shuttle mission lasts 8-12 days, and Chawla hopes to participate in several trips over many years. "This is very far fetched for me," she commented, with a twinkle in her eye. "Yet when I went to college [in India], I had pictures of space shuttles in my dorm -- I'm surprised I even had access to them."

After qualifying for her pilot's license in 1987, Chawla said she began to think seriously of applying to the space shuttle program. Although not a requirement for acceptance, the pilot's license gave her confidence, "because you know you can master other challenges when the time comes." Plus, she has always loved flying.

Flying is `Sheer Fun'

"I like airplanes -- it's that simple. The theoretical side is mentally challenging but flying for me is sheer fun. It appeals to all my senses. The [astronaut] job requires a technical background and a strong desire...to go out in the blue yonder," she said.

Chawla, a U.S. citizen, was born in Karnal, India. She received her B.S. in aeronautical engineering at Punjab Engineering College in 1982, her M.S. in aerospace engineering at the University of Texas in 1984, and her doctorate in aerospace engineering at the University of Colorado, Boulder, in 1988.

Her husband, Jean-Pierre Harrison, is enthusiastic about Chawla's new venture. After relocating to Houston, for what Chawla hopes will be "a very long time," Harrison plans to continue working as a flight instructor and writer. The couple has no children. Chawla said the rest of her family, who reside in India, support her endeavor.

Former Ames scientist Steve Robinson, who has worked at Langley Research Center since 1990, was also selected for the 1995 shuttle program. At Ames, Robinson was one of the principal investigators of turbulent boundary layers, in which he studied the origins of turbulence.

Next Article | Contents | Main Menu | NAS Home

Are Benchmarks the Most Effective Way to Select the `Best' MPP?

by Elisabeth Wechsler

To what extent can benchmarks determine the best MPP (massively parallel processor) performance? That was the question addressed on November 15 at Supercomputing '94 in a session entitled, "Selection of the Best MPP: Do Benchmarks Really Tell the Whole Story?"

While answers varied among the six panelists, most cautioned the Washington, D.C. audience of several hundred against putting too much emphasis on benchmarks as the sole basis for selecting the `best' MPP.

The intent of the panel, chaired by Myron Ginsberg, consultant systems engineer at Electronic Data Systems Advanced Computing Center, was "to provide the audience with useful, practical advice to assess MPP performance, as well as to expose deceptive claims about MPPs."

No Selection Method Is Foolproof

One caveat offered by Ginsberg in his introduction was that "it's impossible at present to provide a simplistic, foolproof method to help you pick the `best' machine to satisfy your specific application needs. And even if that were possible, the rapidly changing user requirements plus the continuously evolving collection of new machines would soon make that selection obsolete."

He went on to say that "simply testing one or more kernels on a new machine is insufficient since slight program modifications can produce a change of as much as two orders of magnitude in performance, and thus could lead a user to erroneous conclusions about the performance of the entire application."

Ginsberg asked the panelists rhetorically: "Can performance with any benchmark suite accurately predict performance for a specific industrial application?" He noted that "this is the dilemma that potential MPP users face in trying to determine if their application should be ported to an MPP and, if so, how to decide which one is best."

Measure Development Cost

Charles Grassl, senior benchmarking analyst at Cray Research Inc., cautioned the audience against benchmark biases: "Running benchmarks `as is' sounds objective, but all tests have potential biases." He would prefer to see the development cost measured, rather than a price/performance ratio, which he believes is a "dependent parameter because it's market driven." For Grassl, "any type of metric summaries are deceptive, because of implicit weighting."

Another opinion was offered by Roger Hockney, independent consultant in parallel computing (formerly associated with Southampton and Reading Universities, UK): "The most useful benchmark is the best-worst performance, which can be tested on a single node with no parallel benchmark required." Best-worst performance refers to which machine is best for the code with the worst performance.

Copyright and Disclosure Constraints

Walter Kohler, senior software engineering manager at Digital Equipment Corp. (DEC), commented that benchmarks for commercial applications provided by vendors are constrained by copyright and nondisclosure laws. He deals with online transaction processing and decision support systems, in which the objective is "to maximize throughput for lots of concurrent users."

As DEC's primary representative to the Transaction Processing Performance Council (TPC), Kohler views cost -- which he defined as the hardware purchase and software maintenance over five years -- as "the most critical factor in selecting the best MPP." For Kohler, "the hardware is the easy part -- the parallel database software is the key," and the "TPC benchmarks aren't code, but a detailed description of what you're allowed to implement."

"Even with one machine, you can get different results from different implementations," noted Aad van der Steen, a consultant in computational mathematics at the University of Utrecht, Netherlands. "There is no real MPP performance because of the vectorization rate, the placement of data in local memory, and local cache behavior."

Structural and Performance Details

"To understand benchmarking, the prerequisites are the details of machine structure, such as processor speed, communications latency, [and] bandwidth; single-node performance; and network performance. You must do some benchmarking in order to understand even these issues. You can't extrapolate to applications regarding their performance on a machine without a proper analysis of these applications -- benchmarks can't answer questions that aren't asked," he said.

Van der Steen advised the audience to "assess the degree of similarity of the user's application to that used in the benchmark" in order to determine a benchmark's relevance.

Los Alamos National Laboratory does benchmarking "to advance the state of computer architecture to serve the future needs of Los Alamos, not to select which machine to purchase," said Harvey Wasserman, staff member of the computer research and applications group. He believes that applications are the most important basis for benchmarks, noting that his view differed from those of other panelists.

Wasserman said the purpose of a benchmark is "not to select a machine that is minimally parallel, but [to] tell you where the bottlenecks are."

He concluded that "MPPs have enormous potential and there's a need for orders of magnitude more than computing power." High sustainable computing rates on real applications is needed, he added. "For easily programmable supercomputers, we need flexible architectures."

Next Article | Contents | Main Menu | NAS Home

Industry Harnesses Clustered Workstations To Squeeze Extra Cycles

by Elisabeth Wechsler

Three case studies of recent aeronautics industry experience show how workstation clusters are being used as an alternative form of high-performance computing. An overview of the NAS workstation cluster project will be covered in the May-June issue of NAS News.

McDonnell Douglas Adapts PVM To Meet CFD Production Code Challenges

On any given night or weekend, McDonnell Douglas has as many as 400 workstations (with an average of 200 per session) divided into clusters of 20 workstations per parallel job doing CFD processing, according to Ray Cosner, CFD Applications Manager at McDonnell Douglas Aerospace, St. Louis.

Using the public domain Parallel Virtual Machine (PVM) software as a foundation, his group implemented and adapted PVM's parallel processing capability into its production CFD code, NASTD, in five months. This McDonnell Douglas proprietary code functions similarly to multi-block structured NASA codes such as OVERFLOW, CFL3D, and NPARC/PARC3D.

Cosner is pleased with the use of PVM-based parallel processing on workstation clusters for large-scale time-critical CFD analyses. "Our experience has met or exceeded our expectations," he said. "By utilizing otherwise wasted workstation computing cycles, we've expanded our useful in-house CFD computing capacity by perhaps 20 times."

A small portion of the jobs run on a heterogeneous collection of workstations, and the remainder use Hewlett Packard models of various types and sizes. However, in Cosner's opinion "heterogeneity with PVM implementations is a non-issue." The real question is one of network topology: "How many network links must be crossed when grouping workstations together to run a job?"

Two years ago, McDonnell Douglas began running selected jobs in parallel on workstation clusters while production jobs were running in single mode. At this point, some problems with fault tolerance were detected. "We believe our attention to this class of issues is one reason our experience has been so positive," Cosner said.

Whereas individual workstations and network elements may be 99 percent fault free for a short job, a cluster of 20 workstations and 50 network elements offer "a very high probability of disruption" for a long-running parallel job, Cosner explained. "A straight installation without considering fault tolerance fails to address a parallel job's need to deal with network hiccups."

After the fault tolerance problem was solved, the workstation clusters were ready for "unrestricted production" in October 1993. Solutions are now checkpointed at user-controlled intervals and various tests are conducted to monitor the health of the network links and the processing nodes.

"If a fault is detected, the parallel system is reconfigured without the offending elements and the job is restarted from the last checkpoint," he said. "We've had jobs successfully reconfigure and restart as many as three times in one night."

These defensive measures have boosted McDonnell Douglas's single-job overnight success rate to about 95 percent. "This reliability is essential for user acceptance of workstation cluster computing," Cosner said. Without these measures, he estimates the success rate for a simple PVM implementation to be well below 50 percent.

Cosner noted that McDonnell Douglas has not yet tried to optimize operating systems, CFD algorithms, database design, or message composition to further increase efficiency in its workstation cluster environment, admitting that "this is an immediate research need."

However, Cosner is convinced that important improvements can still be made. For example, he noted that "the process of periodic solution checkpointing wastes computing cycles." He would like to see "the current 5 percent failure rate reduced substantially, with checkpointing overhead also reduced dramatically," and called for "an improved set of diagnostic tools to improve efficiency and detect failures affecting this type of processing."

Other areas where Cosner would like to see improvement are job scheduling and load balancing. "We've had chronic problems with PVM jobs (which run in the background) stalling because a higher priority process is running on the same node." He explained that this can lead to "all sorts of degradations, including CPU and memory conflicts, and `page thrashing,' which accounts for most of our current overnight job failures."

A related problem is "runaway" or "hung" processes (typically, non-CFD tasks such as system accounting processes), which are not detected until they "effectively kill a PVM job," he said.

The only corporate problem with the workstation cluster project was "convincing the rest of the company that hardware is not personal property." Because most of the competing work is related to interactive computer-aided design (CAD), restricting access to periods between 7:00 p.m. and 6:00 a.m. and on weekends "wasn't much of an imposition," he said.

However, if a CAD engineer came in early, that work would take priority. "This caused the most damaging scenario," Cosner said, adding: "If one of our jobs wouldn't shut down at the end of the night, this also created problems."

Cosner believes that the workstation cluster parallel computing at McDonnell Douglas has been "a highly positive experience, providing supercomputer equivalent cycles in a cost-constrained environment." He hopes "to build on this base for a longer-range evolution into more powerful parallel implementations."

[TOP]

Teamwork Plus 600 Off-duty Workstations Reduce Design Time, Cost at Pratt & Whitney

Pratt & Whitney (P&W) has achieved the throughput equivalent of 16 CRAY C90 CPUs by using networked workstations at two sites, according to Dan Minior, Manager of Aerodynamic and Mechanical Integrity Systems for Management Information Systems.

Computing at P&W is driven by customer requirements to increase performance and the company's desire to reduce design time and cost. P&W's goal is to maximize asset utilization, and Minior believes that loosely clustered workstations contribute to that goal.

By running jobs on parallel workstations during off hours to harvest cycles equivalent to approximately 75-90 CRAY X-MPs, P&W has demonstrated that this approach is a "viable alternative to investing in more costly supercomputers," Minior said. Just three years ago, the same approach delivered only 2-3 CRAY X-MP equivalent cycles, he added.

P&W began a transition from mainframes to workstations at two domestic sites (East Hartford and West Palm Beach) in 1988-89, and achieved its first major success with parallel workstations in 1992, when Sun SPARC 2s were used in a PW4000 fan redesign. "Required accuracy was achieved, lead time was reduced from 18 to 6 months, and savings in excess of $20 million were realized from improved aerodynamic efficiency and reduced testing," Minior said.

In the last six months, P&W has completed the redesign of the high pressure compressor for the advanced version of the PW4000 engine. Efficiency was improved substantially without an adverse effect on stall margin. He predicted that over the life of an aircraft (approximately 20 years) the fleet of P&W's next generation engines "may obtain fuel savings of $1 billion."

To complete the redesign, P&W harnessed 600 workstations in East Hartford plus 300 in West Palm Beach during off hours. During the day, these workstations are used for CAD work, which requires a high degree of interaction with engineers. As first-shift employees leave between 3:30 p.m. and 6:00 p.m., their workstations become part of a pool for calculations until 6:00 a.m., when they revert to daytime use. P&W also uses the workstation network on weekends, holidays--and even during the annual facility shutdown period, Minior said.

Results show that P&W has been able to establish a production environment with workstation clusters, and Minior attributes much of the credit to "a high degree of teamwork." Eleven engineers representing networking, operating support, application developers, and end users provide applications support to all of P&W.

One sub-team spent about a year on reliability issues--addressing a variety of glitches, from problems with UNIX internals to mishaps caused by evening cleaning crews. As a result of this focus, Minior claims that P&W's reliability for its workstation clusters improved from 30 percent to 93-99 percent.

The P&W network consists mostly of Sun workstations of varying sizes and types, one operating system, and a variety of software applications. P&W developed its own communication software--PROWESS (Parallel Running Of Workstations Employing Sockets).

P&W also uses Load Sharing Facility (LSF) software, developed by Platform Computing, Toronto, to maintain "a stable production environment," Minior said. LSF helps with resource monitoring, system management, and job scheduling. Presently, LSF works for serial batch only; however, Minior is working with Platform Computing to modify the software to handle parallel jobs.

Minior said that further work is needed on system software and applications to work with codes using network communications. "We haven't solved all the problems, but we see a great advantage to this alternative form of compute power," he said. "The future should continue to bring us high utilization of equipment and substantial achievement in compute power, at low cost."

The next steps, as Minior sees it, are improved simulation software, faster workstations, faster networks, and improved message-passing. "This will facilitate workstation clustering, and will enable NASA and industry to work together to advance the technology, accelerate the pace of deployment, and improve U.S. competitiveness," he said, adding that he appreciates the support he has received "from all the NASA sites, particularly Lewis Research Center."

[TOP]

'Deployable Software' Is the Key at Boeing

Tom Wicks, Program Manager of High Performance Computing at Boeing Computer Services, stated that his company's investigation of cluster computing started with an attempt to "run parallel and sequential production codes in a transparent manner." The biggest challenge this effort faces is "deployable software" that allows clustered workstations to do message passing and job scheduling.

Boeing's work, which began in 1992, has focused on the issue of "when is the right time and what are the right applications to use." Wicks believes that cluster computing is "applications dependent" and that not every application is appropriate for this environment.

Boeing has been conducting a variety of cluster pilot tests, including interdepartmental testing on two workstation clusters. One cluster, completed last year, involved 14-20 workstations dedicated to propulsion applications. The other cluster, involving 14 workstations for helicopter applications, is just beginning. Some of the codes have extensive communications requirements and others don't. Boeing's efforts have also addressed application codes involving structures, propulsion, aerodynamics, and electromagnetics.

"The idea is to run some applications on mixed clusters," he said. "For example, 15-20 workstations running structures' applications may not have enough cycles, so we may want to run these applications on another group's cluster across the country. In this situation, we want software with transparency, not only to our department but to other work groups as well," Wicks said.

The Boeing tests address network communications issues as both a resource and a requirement for job scheduling purposes, Wicks said. The company is exploring a control system, using fuzzy logic--a mathematical technique for dealing with imprecise data that is more analogous to human logic--which prioritizes the allocation of parallel jobs to, and their suspension from, a cluster of networked workstations. The main objective of this work is to facilitate the harnessing of computational resources from clusters of unused machines and networks, in order to run sequential and parallel jobs, he said.

"If engineers are to rely on workstation clusters, then reliable job scheduling that ensures fault tolerance, interoperability of systems software, and load balancing is a necessary requirement. At present, most job scheduling software available for the cluster environment cannot schedule a job based on network load," Wicks said, adding that "most of the job scheduling software takes into account the status of the CPUs, not the network."

"We'd like to submit the code through some mechanism to a job scheduler that will assign the most appropriate workstation cluster based on communications, CPU, and disk requirements," The work may be done sequentially on the person's desk, a shared memory machine, or maybe a workstation cluster in another department."

Wicks said Boeing has tried most of the major hardware vendors for workstations and supercomputers, and wants the whole environment to be considered for deployment. "We've gone beyond proof of principle to: How do you deploy in a production paradigm?" He noted that issues such as system administration, security, network configuration, infrastructure, and integration issues still need to be resolved.

Two years ago, a typical attitude at Boeing was, "my machine is my machine," Wicks said. "Now the bigger problem is reliability of workstation clusters: If you can't do your work because the software isn't working properly, you'll be reluctant to participate in a cluster project."

In conclusion, Wicks said, "We are cautious about the potential of workstation clusters. We're concerned about putting a production paradigm in place before the right software and support infrastructure is there to handle it, and reliable, commercial software has not been as available as we'd like. We want off-the-shelf software products to ensure that they're properly supported."

[TOP]

Next Article | Contents | Main Menu | NAS Home

Next Article | Contents | Main Menu | NAS Home

Vectorizing Linear Recurrences

by R.K. Owen

A technique called "recursive doubling" offers a valuable way to help optimize first order linear recurrences by casting the linear recurrence into a vectorizable loop. Although the number of operations increases with this technique, it is offset by the use of vectorization. On the CRAY C90, simple forms of linear recurrences can give substantial performance advantages.

First order linear recurrences of the form $\begin{array}{c} x_1 = b_1 \\ x_i = a_i x_{i-1} + b_i \end{array}$ where *i*= *2*,..., *n* are not uncommon in numerical calculations. Such linear recurrences can be used to evaluate the gamma function $\Gamma(z+1) = z\Gamma(z)$ and related functions, the Bernoulli and Euler polynomials, and (a more typical example) polynomial evaluation of a truncated power series using Horner's method:

$$\sum_{i=0}^{n} c_i z^i = (\ldots((c_n z + c_{n-1}) z + c_{n-2}) \ldots) z + c_0$$

Horner's method prevents the inefficient calculation of numerous powers by a transformation to a simple series of multiplications and additions. However, because each element of the linear recursion depends on the previous element, it introduces data dependencies that inhibit vectorization. Recursive doubling can vectorize such recurrences.

A simple example is to fill a vector with successive powers of a given value ($x_i = a^i$). The simple and straightforward scalar code is:

```
      x(1) = a
CDIR$ NEXTSCALAR
      do 100 j = 2, maxpow
        x(j) = x(j-1)*a
  100 continue
```

The comment line starting with "CDIR$" is a Cray compiler directive that tells the compiler to avoid heroic efforts to vectorize the next loop. This decreases CPU time by a factor of two for the CF77 compiler system by avoiding needless runtime conditionals for a loop that can't possibly vectorize. The same task performed using recursive doubling is implemented as follows:

```
      x(1) = a
```

```
        j = 1
   10   continue
CDIR$ IVDEP
        do 100 k = j+1, min(2*j,maxpow)
          x(k) = x(j)*x(k-j)
  100   continue
        j = 2*j
        if (j .lt. maxpow) goto 10
```

The directive "IVDEP" tells the compiler to vectorize the next loop and ignore vector dependencies. The following diagram helps to visualize what's happening:

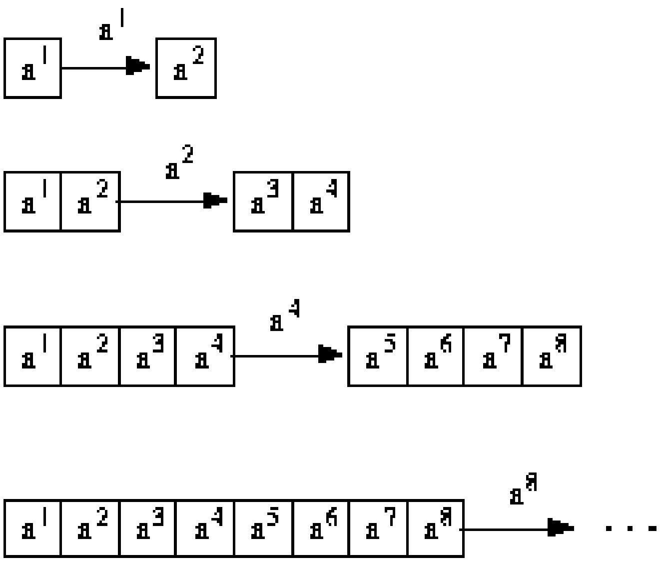

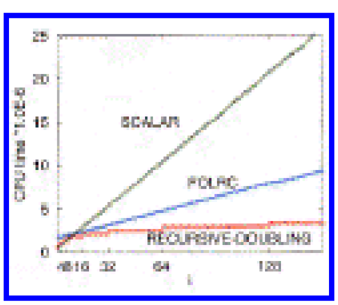

Each run through the vector loop is doubling the length of the vector and only refers to data not within that part of the vector being evaluated -- therefore, no vector data dependencies. The CPU timing results (*see* *Figure 1*) compare the above scalar and vectorizing recursive doubling code with the Cray routine, FOLRC. The recursive doubling algorithm gives CPU times that scale *O*(logn), as opposed to scaling linearly, as with the scalar and FOLRC case, even though the total number of operations increase as *O*(logn) the number of vector operations increases only as *O*(logn) for recursive doubling.

Fortran90 Yields `Surprising' Results

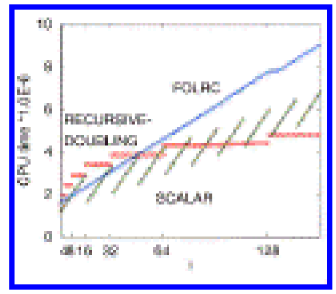

Repeating the above exercise in Fortran90 gives a surprisingly different outcome (*see* Figure 2). The scalar Fortran90 code looks approximately the same, and the entire recursive doubling routine is shown here:

```
SUBROUTINE vlinrec(n,a,x)
        integer, intent(IN) :: n
        real, intent(IN) :: a
        real, dimension(n),intent(OUT) :: x
        integer :: i,k, min
        intrinsic min

        x(1) = a
        k = 1
         do                            ! an infinite loop
                if (k >= n) EXIT     ! do while k < n
!DIR$ IVDEP
                do j = k + 1, min(2*k, n)
                        x(j) = x(k) * x(j-k)
                enddo
                k = 2*k
        enddo
        return
END SUBROUTINE vlinrec
```

The Cray Fortran90 compiler does a tremendous job of optimizing code, particularly scalar loops. The recursive doubling code doesn't surpass the scalar loop in performance until the vector length is approximately 100, as opposed to about 10 for the Fortran77 code. The Cray f90 compiler unrolls the scalar loop 16 times, which accounts for the saw-tooth pattern in the scalar CPU timings.

A practical example of using the above code is to fill a vector with a table of cosines, $t_k = \cos \frac{2\pi}{360} k$ where $k = 0, \ldots, 359$, using the mathematical identity $\cos(nx) = \Re[(e^{ix})^n]$.

Table 1 (*below*) compares the scalar and recursive doubling loop with the direct evaluation of the cosine table with the Fortran intrinsic, COS( ). The speed-up gained by using recursive doubling only becomes significant if such tables need to be computed repeatedly.

Table 1: CPU time (*1.0E-7)

```
------------------------------
            cf77  f90
```

Intrinsic	331	342
Scalar	311	492
Recursive Doubling	148	157

The full first-order linear recursion described at the beginning of this article is a little more difficult to implement using recursive doubling, and leads to the following Fortran90 code:

```
        x(1:n) = b(1:n)
        k = 1
        do
                if (k >= n) EXIT    ! do
                    while k < n
                ! array syntax
                x(1+k:n) = a(1+k:n) *
                    x(1:n-k) + x(1+k:n)
                a(1+k:n) = a(1+k:n) *
                    a(1:n-k)
                k = 2*k
        enddo
```

Note that two vectorized loops are implemented within the body of the doubling loop. This is indicated by the array syntax notation. The overhead of setting up two vector loops outweighs the possible CPU savings. In this full implementation of the first order linear recurrence, the highly optimized scalar loop out-performs the recursive doubling method. In this case the Cray routine FOLRP gives the best performance (*see* *Figure 3*).

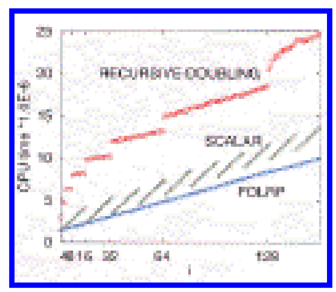

For more information on using recursive doubling, send email to **rkowen@nas.nasa.gov**.

R.K. Owen received his Ph.D. in physics from the University of California, Berekely in 1990. As part of the NAS scientific consulting team, he specializes in math libraries, numberical methods, and object-oriented design and programming. Other areas of expertse include Fortran 77 and 90, C, and C++.

Next Article | Contents | Main Menu

NAS Home

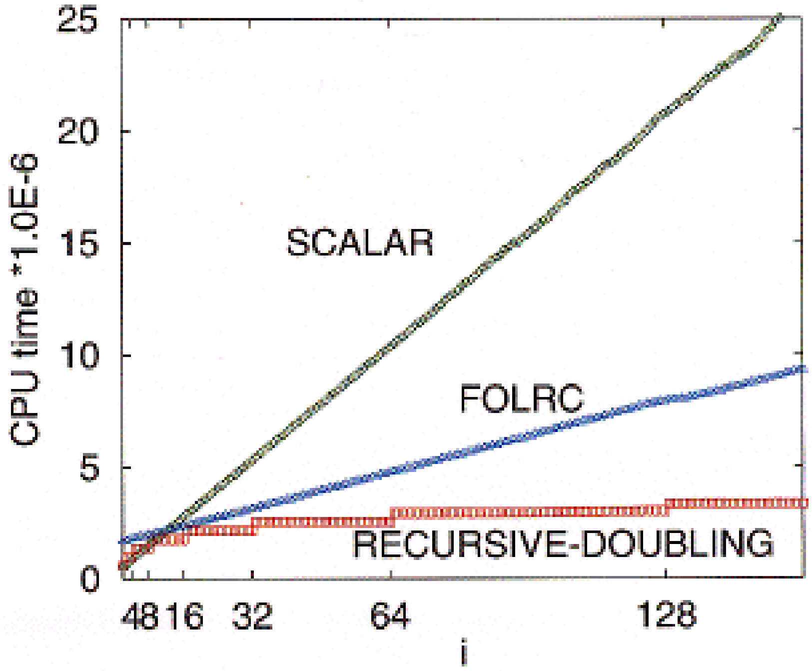
CPU time *1.0E-6
0
5
10
15
20
25
SCALAR
FOLRC
RECURSIVE-DOUBLING
4
8
16
32
64
128
i

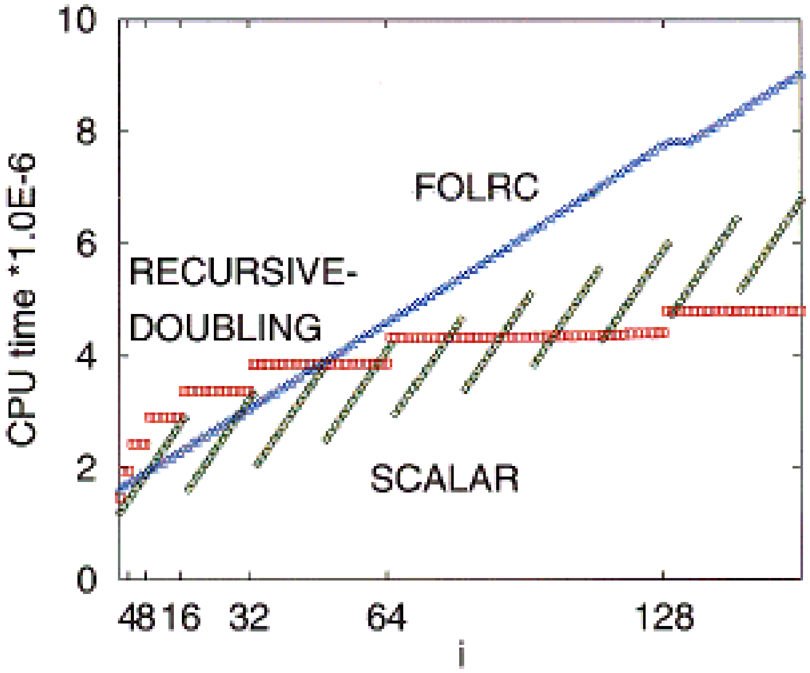
10
8
6
4
2
0
CPU time *1.0E-6
FOLRC
RECURSIVE-
DOUBLING
SCALAR
4816
32
64
128
i

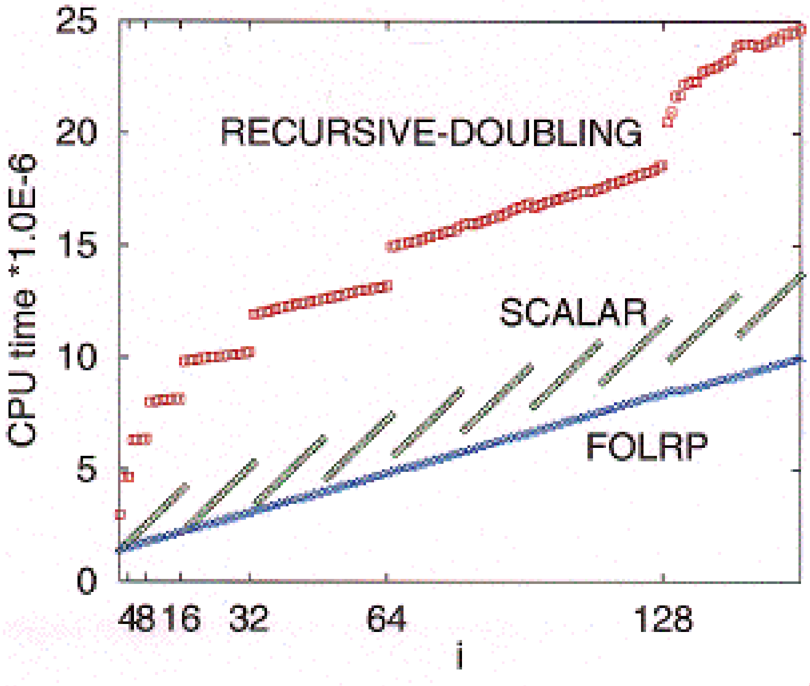
25
20
15
10
5
0
CPU time *1.0E-6
RECURSIVE-DOUBLING
SCALAR
FOLRP
4816
32
64
128
i

No Lack of Challenges For Gigabit Networks

by Elisabeth Wechsler

John Lekashman, manager of networks and mass storage at NAS, joined six other speakers at "Gigabit Networking: State of the Art and Applications," a roundtable session held at Supercomputing '94 on November 17 in Washington, D.C. Speakers offered candid comments about using different networks at their organizations' sites. About 200 conference participants attended.

Moderator Bill Boas, of Essential Communications, said in his introduction that the roundtable's goal was "to share technology, product, implementation, and operation experience with people who have installed and operated gigabit networks at the conference and at their home sites."

Wide Range of Organizations

In addition to Lekashman, the speakers represented a wide range of organizations, including Ford Motor Co., National Center for Supercomputing Applications (NCSA) at the University of Illinois, Sandia National Laboratory, University of Pennsylvania, and Los Alamos National Laboratory.

Speakers surveyed a number of "commercially ready" standard interconnects, protocols, and architectures for gigabit networking, such as ATM (Asynchronous Transfer Mode), SONET (Synchronous Optical Network), FCS (Fibre Channel System), and HiPPI (High Performance Parallel Interface). Their presentations revealed a wide range of experiences with gigabit networking. Some of the applications discussed included cluster computing, storage management, and digital media distribution--all in the context of the speakers' organization environments.

Anticipated Needs at NAS

Lekashman listed the anticipated requirements of gigabit networking at the NAS Facility as:

- higher aggregate performance
- higher single-transfer performance
- 400+ megabytes (MB) per second data transfer speed needed for new applications, such as the NAS Virtual Wind Tunnel
- more reliability

- use of standard protocols--TCP/IP
- wide vendor support

Of these, the most important are the use of standard protocols and wide vendor support, he said.

HiPPI Extension to Desktop

Lekashman predicts that over the next two years HiPPI network technology will be extended to the desktop and that development will be highly focused on ATM and Fibre Channels.

On Lekashman's wish list for gigabit networks are the enabling of 500 MB per second single-data transfers and the disappearance of distinctions between local area networks (LAN), metropolitan area networks (MAN), and wide area networks (WAN).

Lekashman believes that latency (that is, the communication time between sending and receiving data) will "still be an issue" in two years. He lamented that "vendors make the same mistakes over and over." For example, he said that vendors "start from scratch on the problem of delivered network throughput with every operating system upgrade."

Don Tolmie, of Los Alamos National Laboratory (LANL), commented on the status of develop-ment extensions to HiPPI. These include operating serial HiPPI over local fiber channel, developing HiPPI-to-SONET gateways, and defining HiPPI-to -ATM compatibility.

A `Fire Hose' for Big Things

HiPPI, which was developed at LANL, is "not good at small things such as multiplexing," Tolmie admitted, but added that "it's a fire hose for
big things." Among its other "shortcomings," Tolmie pointed out that HiPPI is not a mass marketed item, doesn't support lower speeds, doesn't do a good job of multispeed on one network, doesn't perform isochronous communications, and can function well only for limited distances on copper wire.

Tolmie believes that HiPPI is best used, for example, with supercomputers, high end workstations and workstation clusters, RAID storage, RAM disks, and high speed helical scan tape systems.

Ray Cline, of Sandia National Laboratory at Livermore, said that his organization's goal was to "tie together all resources to solve problems, from desktop to high-end computers, and across the country." He added, "We need to deal with what's available."

No Single Solution

Randy Butler, of NCSA at the University of Illinois, said there is "no single solution," nor any "near-term horizons that work," in trying to create gigabit networks. However, he sees a "push for HiPPI."

"Serial HiPPI is making big progress with copper wire for Silicon Graphics Inc. (SGI) and other high-end workstations." NCSA uses FDDI (Fiber Distributed Data Interface) and HiPPI together, and have had a LAN testbed for ATM in process for a year, he said.

Butler thinks the real issue is in host-stack implementation. "We're seeing 98 to 80 percent of extremely expensive processors [dedicated] to achieving acceptable bandwidth. How do you achieve this without spending $10 million on networks?" he asked, joining Lekashman in expressing concern for vendors' tendency to reinvent the wheel with each operating system upgrade.

Steve Gossage, of Sandia National Laboratory at Albuquerque, advised the user audience to "leverage standards whenever possible." His preferred network architecture is ATM, especially the "simpler extensions."

ATM: Missing Pieces

Among the pieces he believes are "still missing" are: a protocol aimed at the gigabit networking environment; low latency throughout hardware and software; low-bit error RAID; and a large-capacity ATM switch (2.5 gigabits) for the entertainment industry, video, and voice.

Gossage thinks that the SONET infrastructure needs to be upgraded to 1 gigabit--although 10 gigabits per second "would be better." Also, he said, the cost of bandwidth needs to go down. "The entertainment industry may help bring prices down by building up infrastructure" and thus significantly increase demand.

Jonathan Smith, of the University of Pennsylvania, talked about the experiences of a consortium his university has joined, which focuses on networking issues. Other members include regional phone companies in the northeastern U.S., IBM, and Massachusetts Institute of Technology.

The consortium has singled out the networking problems of workstations because it believes that workstations will be "the dominant architecture" of the future, Smith said. "The host workstation is where the real work is done [and] the real issue is I/O architecture," he said, noting that supercomputers have "an edge with this."

Network Overhead--A Workstation Problem

"The network overhead doesn't leave much room for workstations to do any work," Smith said, because of the speed-of-light problem. For example, it takes 30 milliseconds to send communications across the country or to parallel computers on the computer room floor. "You can get supercomputer performance from commodity-level workstations with the right network," he added.

Another problem is latency, according to Smith. "The next generation of research must be in hardware and software to reduce latency." He noted that ATM network latency is lower by a factor of 50 percent, and predicts that "if ATM can get that right, it will prevail" as the preferred gigabit network architecture.

Smith also commented that "Ethernet will be around for awhile but eventually will go away in three to five years, along with other `legacy' networks."

In conclusion, Smith observed: "We can't expect vendors to increase memory bandwidth [unless] we can convince them there's enough of a market to make a profit--like the graphics industry did--at least at the workstation level."

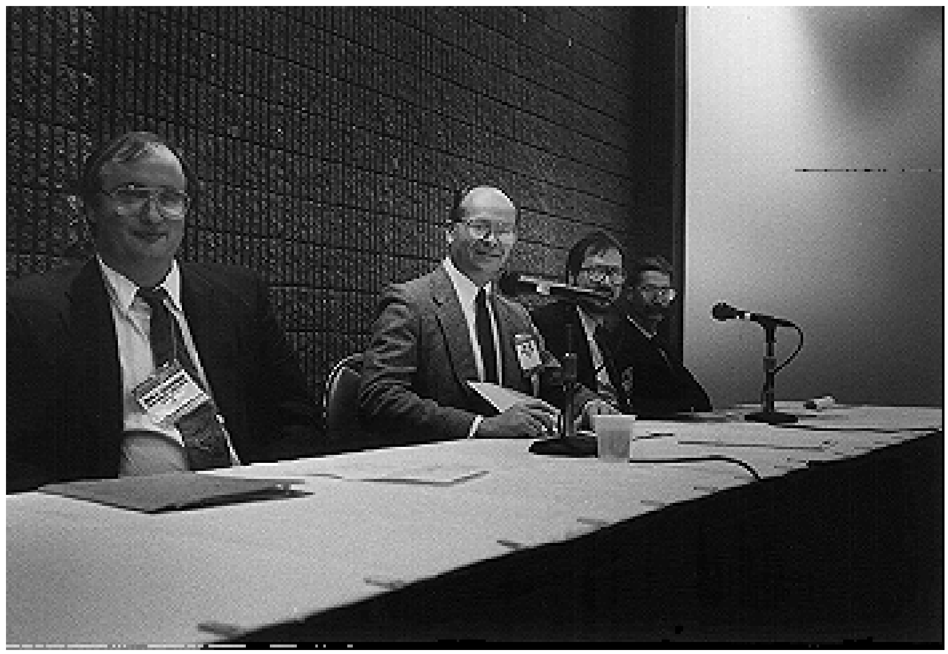

(From left) John Lekashman (NAS), Don Tolmie (Los Alamos National Laboratory), and Ray Cline and Steve Gossage (both of Sandia National Laboratory) formed part of the "Gigabit Networking" panel at Supercomputing '94. Speakers addressed a broad range of issues, including arch itectures, cluster computing, storage management, and digital media distribution, as they affect gigabit networks.

Next Article

Contents | Main Menu | NAS Home

'FASTexpeditions' Lets Users Share, Collaborate on 3D Scientific Data

by Jill Dunbar

"FASTexpeditions" -- an innovative software tool for scientists to share datasets and collaborate on solutions using a World Wide Web (WWW) browser and FAST (Flow Analysis Software Tool-kit) -- is now available on the NAS WWW server. The prototype is already being used for collaborative analysis of CFD solutions among researchers at NASA Ames Research Center, McDonnell Douglas Aerospace, and Stanford University.

Conceived by Val Watson and developed by Jean Clucas, with assistance from John West (all of the NAS Data Analysis Branch), FASTexpeditions gives researchers access to a scientific dataset through a WWW browser, such as Mosaic, and automatically runs an interactive, guided exploration of the dataset. Users at different sites can simultaneously view data on their workstations, and can pass control of the FASTexpedition to one another.

Sophisticated Tools -- But Easy Access

What separates this method from others is that, instead of sending data as pixels to a movie player, data is sent in a form that can be analyzed by sophisticated tools running on a student's or researcher's computer. With movies, up to several megabytes of data can be required to transfer a single animation. A similar or smaller number of bytes is sufficient to transmit a FASTexpeditions dataset -- plus many animations -- over the Internet.

The FASTexpedition tools allow FAST to be driven by Mosaic hyperlinks: By clicking on a link in the Mosaic window, a dataset will be retrieved from a remote site. FAST is started automatically, and the dataset is passed to it. Clicking on additional hyperlinks start FAST animations, which display features of the data the data. These links and explanations in the window guide users through the data -- and, since they are interacting with an actual dataset, users can further explore aspects of the data on their own.

Collaboration on New Techniques

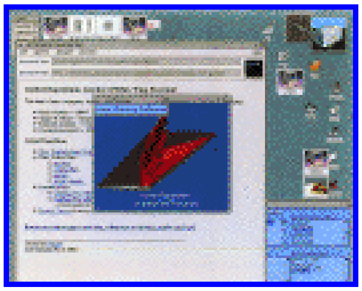

The collaboration with McDonnell Douglas and Stanford, currently being led by the NASA Ames Design Cycles Technologies Branch, uses FASTexpeditions for remote sharing and analysis of CFD solutions. Researchers who are creating new techniques to extract features from CFD solutions link with other researchers who are developing an artificial intelligence (AI) system. These two groups jointly analyze whether particular features meet necessary criteria, then cast the analysis into rules to be used by the AI system. "We have to have remote visual analysis to do this," said Watson. "FASTexpeditions are the way we get the information to each other. Stanford can put its features into a guided expedition, everyone can load it simultaneously, and then we can do joint `what-if' analysis."

'Great Opportunity' for Education

Watson believes that methods such as FASTexpeditions are a great opportunity for the government to contribute to science education because they provide truly three-dimensional (3D), interactive exploration of scientific data using easy-to-learn tools, and allow students to do what-if analysis, as well as facilitate remote collaboration.

Watson has submitted a proposal to bring "scientific expeditions" into schools to the computer graphics and communications communities via the WWW, and has received some important endorsements, including: the board of directors of the National Computer Graphics Association, SIGGRAPH, and the American Institute of Aeronautics and Astronautics Interactive Computer Graphics Technical Committee.

For sites with Internet connections and a Silicon Graphics Inc. workstation, the requirements to run Fastexpeditions are relatively simple: a WWW browser, FAST software, and the FASTexpedition tools. The tools are available at no cost on the NAS WWW server, and FAST can be purchased through COSMIC, Athens, GA. (For information, send email to **service@cosmic.uga.edu**).

For further information, send email to **watson@nas.nasa.gov**. Further details are available on the FASTexpeditions page.

Next Article | Contents | Main Menu | NAS Home

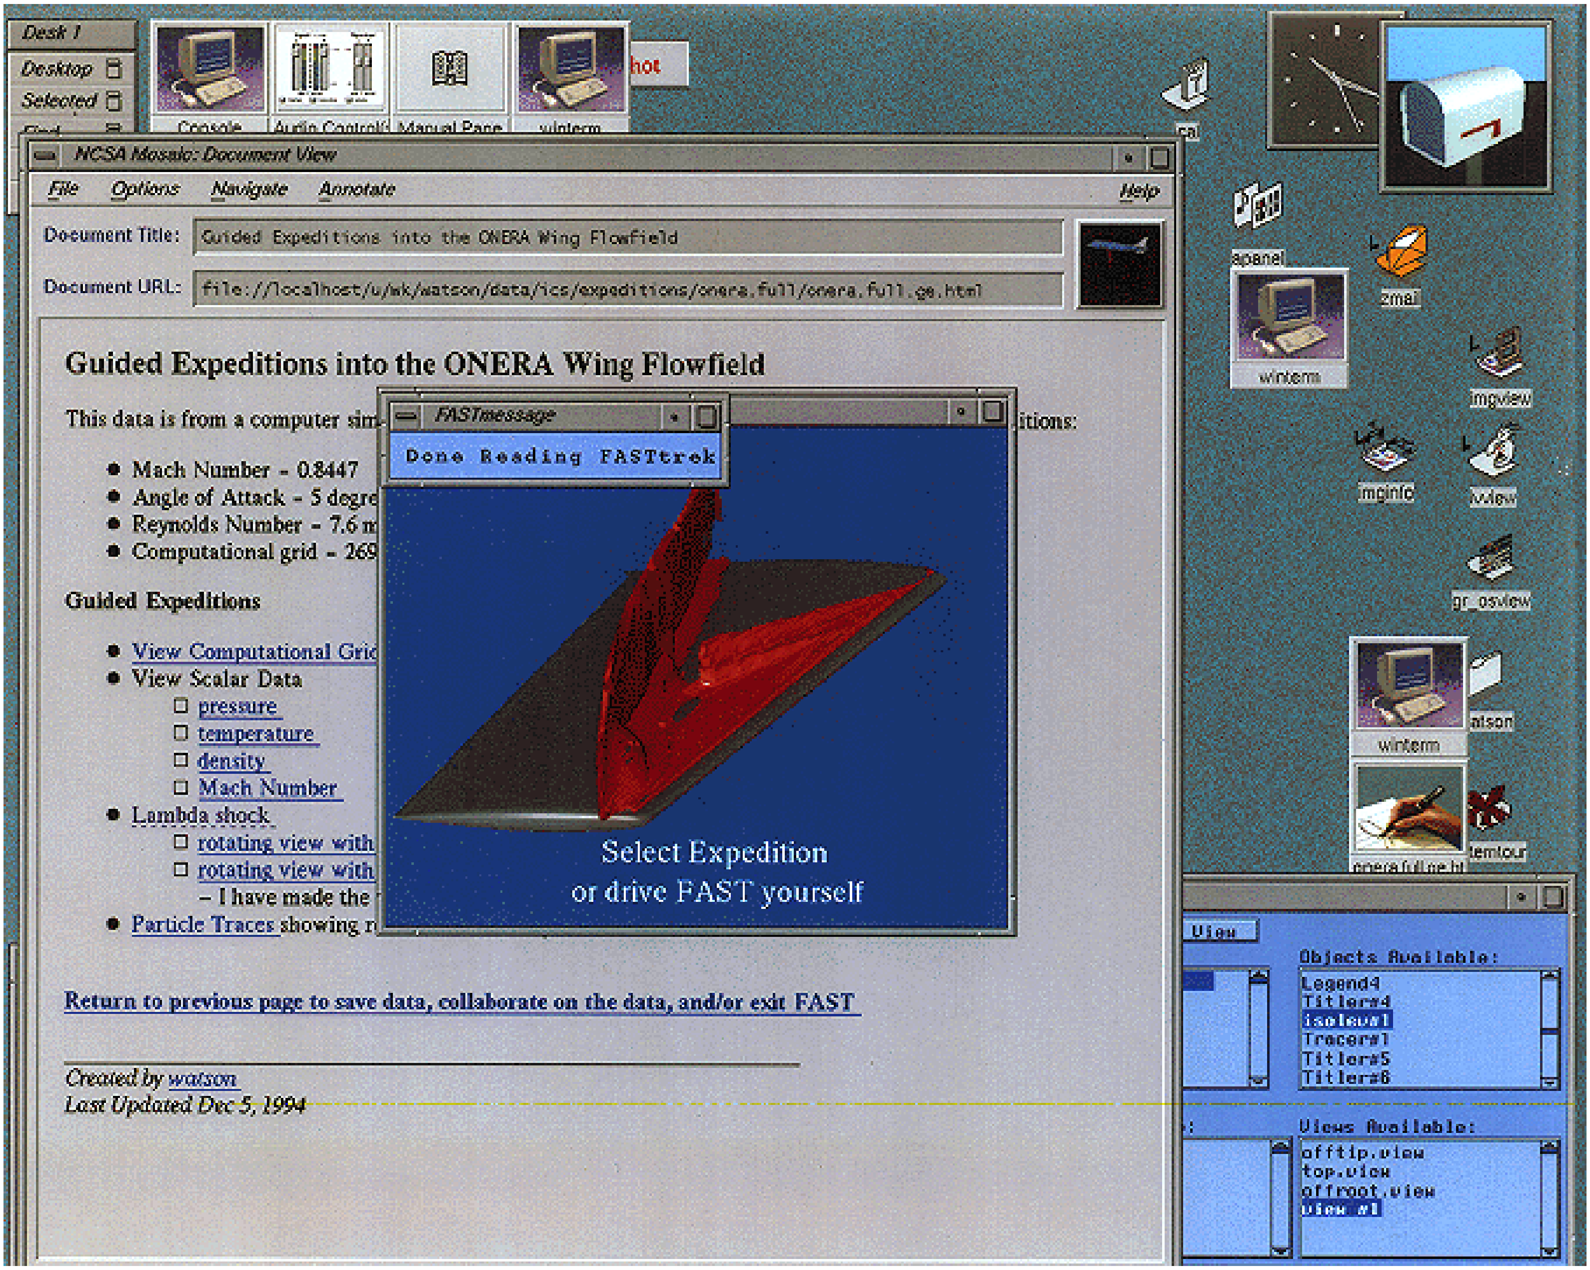

A view of the FASTexpeditions user interface through the Mosaic browser. The inteface allows scientists to do joint visual analysis of data via hyperlinks. The three-dimensional visualization depicts a Lambda shock on a wing, part of a Stanford University dataset being analyzed in collaboration with researchers at NASA Ames Research Center and McDonnell Douglas Aerospace.

to the article

CAVE Focuses on Giving Researchers Real-time Tools

by Marcia Redmond

The CAVE Automatic Virtual Environment (CAVE), a projection-based virtual reality system in which stereoscopic computer graphics create an illusion of immersion, was the focus of a recent NAS New Technology Seminar. Carolina Cruz-Neira, a researcher at Electronic Visualization Laboratory (EVL) at the University of Illinois, Chicago, discussed the progression of CAVE, along with its current capabilities, and her experiences introducing scientists around the country to this technology.

The CAVE "theater" is a 10-ft. cube composed of three rear-projection screens for walls and a down-projection screen for the floor. Added to this are a head-and-hand tracking system and a sound system, which give users a sense of being surrounded (similar to the image shown below).

A researcher in the CAVE virtual reality theater uses a hand-held device to manipulate the Sierpinski tetrahedron. The CAVE environment is controlled by an eight-processor Silicon Graphics Inc. Onyx workstation and three Reality Engines2.

Multiple users can occupy the CAVE theater, allowing scientists to discuss the data in this virtual world and to take turns manipulating objects with a three-button wand-like device.

The technology has been implemented by Cruz-Neira, currently a Ph.D. student at the university. Her current research is focused on providing researchers with tools to dynamically steer scientific simulations in real time by integrating the CAVE environment with massive and scalable parallel computers via high-speed networks.

More information on the CAVE technology is available online from EVL HPCCV Publications.

To contact Cruz-Neira, send email to **carolina@evl.eecs.uic.edu**. To order a copy of the NAS New Technology Seminar videotape, send email to **doc-center@nas.nasa.gov**.

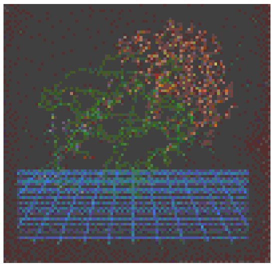

Visualization of the molecular dynamics of cancer; AP21 molecule interacting with water molecules in its vicinity.

Mar - Apr 1995, Vol. 2, No. 9

Executive Editor: Marisa Chancellor

Editor: Jill Dunbar

Senior Writer: Elisabeth Wechsler

Contributing Writers: Kristina D. Miceli, R.K. Owen, Marcia Redmond

Image Enhancements: Chris Gong

Other Contributors: David Bailey, David Barkai, Eric Barszcz, Kalpana Chawla, Jean Clucas, Ray Cosner, Catherine Cruz-Neira, Al Globus, Jonathan Hahn, Toby Harness, Eric Hibbard, Ed Hook, John Lekashman, Bob Meakin, Dan Minior, Eugene Miya, Chuck Niggley, Lew Schiff, Cathy Schulbach, Val Watson, Sisira Weeratunga, Tom Wicks, David Yip

Editorial Board: Marisa Chancellor, Jill Dunbar, Chris Gong, Kristina Miceli, Serge Polevitzky, Pamela Walatka, Elisabeth Wechsler, Rita Williams

Special Thanks: CAVE virtual reality theater images courtesy of Carolina Cruz-Neira, Electronic Visualization Laboratory (EVL) and Paul Bash, Argonne National Laboratory; AP21 molecule courtesy of Randy Hudson and Alan Verlo, EVL and Jim Crutchfield, Sante Fe Institute.

NAS NEWS

Volume 2, Number 9 March – April 1995

160-node SP2 'Scales Well'

by Elisabeth Wechsler

The 160-node IBM SP2 at NAS performed an average of 1.4 times faster than the 16-processor NAS CRAY C90 on three pseudo-application NAS Parallel Benchmarks (NPB) run on January 17. According to Toby Harness, of the NAS parallel systems group, "All three benchmarks did better than expected, which shows that the codes are scaling well," based on test results for the former 64-node and 128-node SP2 configurations. The tests were run on full-production NAS machines by NAS staff, Harness noted.

A full report on the NPB, including these latest results, is available on the World Wide Web at: http://www.nas.nasa.gov/NAS/Technical/NPB/NPBindex.html.

The 160-node configuration, which includes the final 32 nodes specified under the NASA Computational Aerosciences-funded Cooperative Research Agreement, allocates 144 nodes for scientific jobs during off hours (or 128 nodes during the day), reserving 16 nodes for system software development. (More details on the 160-node configuration are available from the World Wide Web at: http://lovelace.nas.nasa.gov/Parallel/SP2/nas_sp2_home.html.)

Other improvements to the SP2 include upgraded "general availability" IBM system software; six HIPPI channels in full production, allowing the SP2 to access NAS high-performance systems over the HIPPI network; and support for P-OVERFLOW, completed ahead of schedule.

P-OVERFLOW, a subset of OVERFLOW's most frequently used functionalities, is intended for distributed memory MIMD (multiple instruction, multiple data) computers, according to Sisira Weeratunga, who ported the code to the SP2. OVERFLOW is a general-purpose implicit flow solver, developed by Pieter Buning at Ames Research Center. It uses the concept of a composite of logically structured overset grids to solve flows around complex geometries.

The performance of P-OVERFLOW on the 160-node SP2 is "good," Weeratunga said, adding that "fewer than 8 nodes of the SP2 achieve the performance of 1 CPU of a C90."

Performance results comparing three message-passing libraries—MPL, MPI, and PVMe—for the

Continued on page 7

THIS ISSUE

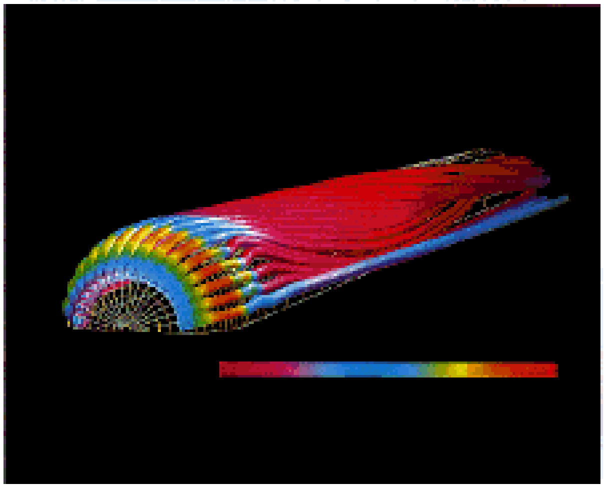

The reversible momentum flux density tensor depicted as hyperstreamlines in a flow past a hemisphere cylinder. The trajectories of the hyperstreamlines are streamlines of the velocity field. The eigenvalues used to compute the cross-sectional geometry are degenerate, resulting in a circular cross-section, with the diameter of the section encoded by local pressure. The color map represents the density of kinetic energy.

NAS-supported Visualization Research Important for Studying Complex Tensor Data

by Kristina D. Miceli

Researchers at Stanford University, working under a NAS-sponsored visualization grant, are devising creative techniques to help scientists depict the complex information contained in tensor data. Principal investigator Lambertus Hesselink and graduate student Thierry Delmarcelle focused the first year of their research on hyperstreamlines and tensor field topology.

New Tools for Viewing Tensor Data

Computer-generated visualizations of computed flow fields typically mimic the images obtained from experimental studies. For example, streamlines computed from computational fluid dynamics (CFD) simulations mimic the patterns created when smoke is injected in wind tunnel studies. Techniques for visualizing scalar and vector field data are widely available and used by scientists. In contrast, techniques to visualize tensor fields—fundamental in engineering and physical sciences—are not readily available, making it difficult for scientists to effectively analyze tensor data. Reasons for the lack of visualization techniques for tensor data include the complex, multivariate nature of the data and the fact that no commonly used experimental analogy exists for visualizing tensor data. People are simply not used to creating mental images of tensor data.

Second-order tensor fields, which consist of a 3 x 3 matrix defined at each node in a computational grid, are frequently computed by researchers in computational aerosciences disciplines such as fluid dynamics and structural dynamics. Tensor field variables in CFD include stress, viscous stress, rate of strain, and momentum flux; tensor variables in structural dynamics include stress and strain.

Because of the lack of visualization techniques for analyzing all components of a complex tensor field, scientists typically resort to either simplifying the tensor field data and visualizing it as a scalar field or visualizing the individual vectors that compose the tensor field. Both methods result in the loss of valuable information that might provide new insight into physical problems.

Hyperstreamlines Explained

Vector fields can be visualized effectively using streamlines, which represent the trajectory path of a particle in a flow field. Hyperstreamlines, an extension of the concept of streamlines to second-order tensor fields, provide the scientist with a continuous representation of the tensor field along a three-dimensional path.

Hyperstreamlines are defined by a trajectory path and a cross-sectional geometry at each point in that path. The trajectory path of a hyperstreamline is essentially a streamline computed from one of the eigenvector fields. The cross-sectional geometry of a hyperstreamline is determined by the two remaining eigenvector fields and can either be an ellipse, a circle (if the eigenvectors are "degenerate," or equal to each other), or a cross. Connecting these cross-sectional geometries along the trajectory path results in a hyperstreamline that either looks like a tube (elliptical/circular geometry) or a helix (cross

Continued on page 2